

Predicting Drug Interactions and Mutagenicity with Ensemble Classifiers on Subgraphs of Molecules

Andrew Schaumberg^{1*}, Angela Yu¹, Tatsuhiko Koshi², Xiaochan Zong³, and Santoshkalyan Rayadhurgam⁴

¹*Cornell University and the Tri-Institutional Training Program in Computational Biology and Medicine*

²*Weill Institute for Cell and Molecular Biology and Department of Molecular Biology and Genetics, Cornell University*

³*School of Electrical and Computer Engineering, Cornell University*

⁴*Department of Materials Science and Engineering, Cornell University
Cornell University, Day Hall, Ithaca, NY 14853 USA*

(Dated: Dec 11, 2014)

Abstract

In this study, we intend to solve a mutual information problem in interacting molecules of any type, such as proteins, nucleic acids, and small molecules. Using machine learning techniques, we accurately predict pairwise interactions, which can be of medical and biological importance. Graphs are useful in this problem for their generality to all types of molecules, due to the inherent association of atoms through atomic bonds. Subgraphs can represent different molecular domains. These domains can be biologically significant as most molecules only have portions that are of functional significance and can interact with other domains. Thus, we use subgraphs as features in different machine learning algorithms to predict if two drugs interact and predict potential single molecule effects.

(Wishart et al., 2013) naturally occur in human cells. Function is often characterized by interactions, e.g. KEGG (Kanehisa & Goto, 2000) defines functional biological pathways in terms of pairwise molecular interactions, while DrugBank (Law et al., 2014) defines drug-drug interactions that influence dosage and functional drug-target interactions. Machine learning efficiently finds promising drug candidates, reducing drug discovery cost by only taking drugs to in vitro tests that are predicted to selectively interact with the molecule of interest (Fukunishi, 2009) (Khandelwal et al., 2008).

Problem Definition and Methods

In our first aim, we benchmark our methods on a mutagenicity prediction dataset. In our second aim, we predict the interactions of interest.

Task Definition

Aim 1 Predict mutagenicity based on subgraphs. Mutagenicity is the ability of a compound to change genetic materials. If a drug is mutagenic, it cannot be safely consumed by a patient, even with predicted drug-target interaction.

Aim 2 Predict drug interactions based on their subgraphs. We consider both small molecules and proteins for interactions. An interaction occurs if dosage must be adjusted, due to both drugs being present in a patient (Law et al., 2014).

Algorithms and Methods

Features consist of subgraph counts derived from the dataset. Subgraphs to be included in the feature set depended on parameter choices of height and distance. Height describes the max distance between any two

Introduction

In drug discovery, each new molecular entity costs on average \$1.8 billion (Paul et al., 2010). Finding drug candidates from over 50 million small molecules available in open access databases (Roy et al., 2010), such as PubChem (Bolton et al., 2008) and ChEMBL (Gaulton et al., 2012), have the potential to save both lives and money.

Over 90,000 proteins (González-Porta et al., 2013) and 40,000 small molecule metabolites

Fall 2014 CS5780 Machine Learning final project paper. Update coming with 20,000+ pos examples and bugfix.

*Corresponding author: ajs625@cornell.edu

nodes in the subgraph. Distance is the max number of nodes away two subgraphs are from each other in the full graph.

Rather than use only subgraphs at a single height, we subgraphs for multiple heights up to a maximum height (Figure 1). We let subgraph s at height h among atoms A be defined as:

$$s(h, a \in A, A) = \forall b \in A, d(a, b) \leq h \quad (1)$$

where d is shortest path distance. Let c be the count of occurrences of a subgraph s in a molecule M having atoms A :

$$c(h, A) = \sum_{a \in A} I(s(h, a, A)) \quad (2)$$

where I is the indicator function. We let the feature set F of all counts c at desired heights H fully describe the molecule M having atoms A :

$$F(H, A \in M) = \forall h \in H, \{c_1(h, A), \dots, c_{n=|F|}(h, A)\} \quad (3)$$

For mutagenicity prediction, we considered two kinds of feature sets:

1. Similar to prior published methods (Costa & Grave, 2010), we considered a pair of subgraphs separated by a distance as a feature.
2. More simply, we considered subgraphs at desired heights as features, rather than subgraph pairs.

For drug interaction prediction, our method reduces computational cost in two ways:

1. Consider only a subset of heights up to a maximum height, not all such heights, and
2. do not consider pairs of subgraphs separated by a distance as features.

We also implemented the two kernels from Costa and Grave (2010) and Faulon et al. (2008) and generated kernelized and un-kernelized data (Costa & Grave, 2010) (Faulon et al., 2008). We were unsure as to the possibility of implementing a kernel trick into all of the ensemble classifiers we chose to use, so we had to generate kernelized input data to use in our ensemble classifiers. Thus used three types of datasets:

1. Subgraphs or subgraph pair data.
2. Data after cosine similarity kernel applied, like Faulon 2008.

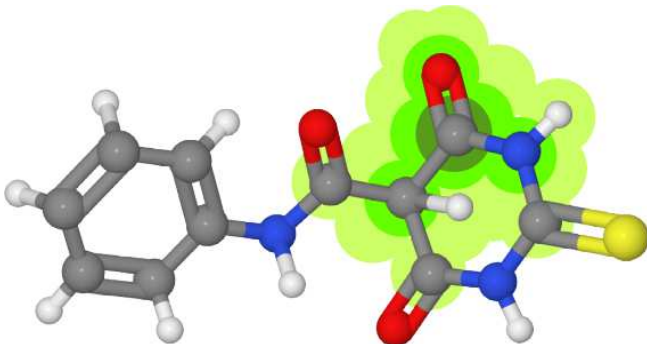


Figure 1. SUBGRAPH HEIGHTS 0, 1, AND 2 IN DARK, LIGHT, AND PALE GREEN REGIONS, RESPECTIVELY

3. Data after NSPDK kernel applied, like Costa and Grave 2010.

We tested convolutional neural network (CNN), random forest (RF), multilayer perceptron (MLP), and support vector machine (SVM) algorithms. Random forest performed best.

Random Forest

Random forests were generated using Scikit-learns RandomForestClassifier (Pedregosa et al., 2011). Random forest subsamples the training data and feature set to build each decision tree in the forest. Each decision tree was allowed to grow fully and both entropy and gini impurity were tried as splitting criteria. 100 decision trees were used to make the random forest. This particular implementation averages the probabilistic prediction of all decision trees in the forest instead of letting each decision tree vote for a classification.

Support Vector Machine

We used SVMlight on all dataset types with all combinations of settings:

1. Linear hyperplane (`svm_learn -t 0`) or radial basis function kernel (`svm_learn -t 2`) classifier.
2. Negative example ratio correction enabled or disabled (`svm_learn -j 0.8`). In the Bursi data, there are 0.8 negative examples for every positive example.
3. Sum of slacks, C , of 1, 10, and 100 (`svm_learn -c 1`).

Table 1. Machine learning algorithm performance: area under ROC and validation set accuracy. Each algorithm’s performance is noted for the top one or two height/distance formats parsed from the Bursi dataset. RF metrics listed as mean \pm stdev. MLP and CNN metrics listed as (min,mean \pm stdev,max). CNN processes each data row as three partitions, each partition being one third of the original row, the first third being subgraphs with the fewest atoms, the last third being subgraphs with the most atoms. This CNN connects the first partition to the first hidden node, the first and second partitions to the second hidden node, the second and third partitions to the third hidden node, and the third partition to the fourth hidden node. For MLP, 3 \rightarrow 1 topology means there are three hidden nodes and one output node. For RF, 100 trees were in the forest, where each tree was allowed to grow out fully. Kernels were not used for any algorithm reported in this table.

ALGO	TOPOLOGY	HEIGHT/DISTANCE	AUROC	VALIDATION ACC	TRIALS
RF	DATA \rightarrow 100TR	H=1	0.89739 \pm 0.024	0.82639 \pm 0.011	100
RF	DATA \rightarrow 100TR	H=[0-3] D=[0-5]	0.89486 \pm 0.007	0.82022 \pm 0.009	100
CNN	3 ^(DATA\rightarrow3) \rightarrow 4 \rightarrow 1	H=[0-3] D=[0-5]	-	(0.78701,0.79085 \pm 0.004,0.79609)	4
MLP	DATA \rightarrow 3 \rightarrow 1	H=2	-	(0.60545,0.76383 \pm 0.044,0.81564)	40
MLP	DATA \rightarrow 3 \rightarrow 1	H=[0-3] D=[0-5]	-	(0.55307,0.72895 \pm 0.108,0.82332)	7

Multilayer Perceptron

We implemented a multilayer perceptron (Mitchell, 1997):

1. One output node and three hidden nodes.
2. Weight updates through backpropagation and momentum.
3. Voted and unvoted perceptrons.
4. Annealed learning rate.
5. Convergence when validation error did not improve after one full iteration through data.
6. For online learning, not batch learning.

Convolutional Neural Network

We implemented a convolutional neural network (LeCun et al., 1998):

1. One output node and four hidden nodes.
2. Data split into thirds by subgraph mass, e.g. when sorting a rows X subgraphs by lowest to highest mass, the $\frac{1}{3}$ subgraphs are grouped as the lightest third, the next $\frac{1}{3}$ subgraphs are grouped into the middle third, and the next $\frac{1}{3}$ subgraphs are grouped into the heaviest third.
3. Other settings identical to MLP.

Experimental Evaluation

Methodology

To allow for direct comparison to prior methods (Costa & Grave, 2010) (Faulon et al., 2008), we opti-

mized for area under the receiver operator curve (AUROC) instead of accuracy. AUROC is a measure of the effect of changing the discrimination threshold of a binary classifier. It is calculated by plotting the true positive rate against the false positive rate over many different threshold values, and then calculate the area under said curve. AUROC and accuracy are different because accuracy can be thought of as a measurement at one threshold value. For completeness, we report training and validation accuracy. We used both 10-fold cross-validation and 100 iterations of randomly splitting the dataset such that $\frac{2}{3}$ makes up a training set and $\frac{1}{3}$ makes up a validation set. We also compare different methods using two-sample un-pooled t-tests with a significance threshold (alpha) of 0.05.

MUTAGENICITY PREDICTION

We computed features at a variety of feature heights (1-12) and distances (0-20). Algorithms we tested include RF, MLP, CNN, and SVM. Statistical significances were determined using un-pooled two-sample t-tests.

We wrote a Bursi data parser that processes the data in two ways: (1) decomposing a molecule into subgraphs (Faulon et al., 2008), and (2) decomposing a molecule into subgraph pairs separated at a distance (Costa & Grave, 2010). Though in prior work (Costa & Grave, 2010), we forwent PTC and NTP HIV analyses for two reasons: (1) their categorical classifications are more complex than the binary classifications in Bursi, and (2) our motivation to focus on interaction prediction for the public good.

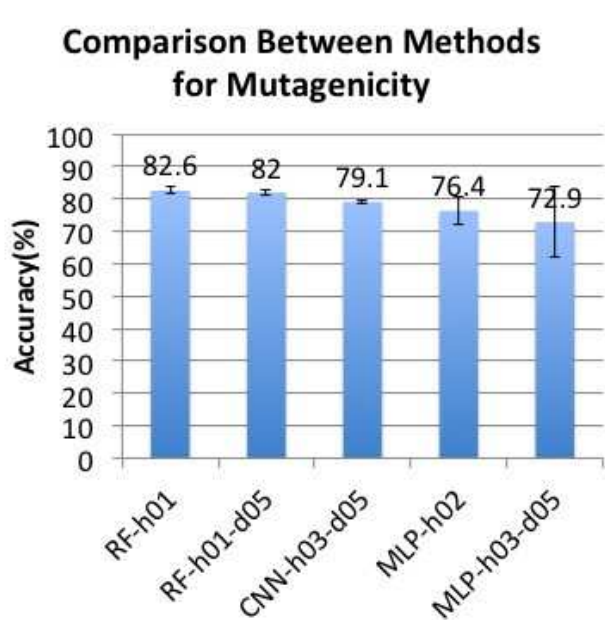


Figure 2. COMPARISON BETWEEN METHODS FOR MUTAGENICITY

INTERACTION PREDICTION

Means and standard deviations of AUROC and accuracies were established using 10-fold cross-validation.

We wrote a DrugBank data parser that decomposes a molecule into subgraphs. For drugs or targets listed in DrugBank where SMILES or protein sequence was not provided, the parser fetches the missing information from PubChem or UniProt, respectively. PubChem results where the chemical name does not match the DrugBank-provided name are rejected. We took two samples from DrugBank for evaluation: (1) 37 positive training examples and 39 negative training examples, and (2) 131 positive training examples and 140 negative training examples.

Results

Among all the algorithms that we tried, random forest had the best result out of the ensemble classifiers (Table 1) (Figure 2). Random forest did significantly better than the other ensemble classifiers (from two-sample unpooled t-tests). A more detailed analysis of the random forest method showed that using height one gave us the best results (Figure 4). Also the two different splitting criteria did not produce any significant changes (from two-sample unpooled t-tests). We compared random forest and previously published methods for mutagenicity prediction, and found that

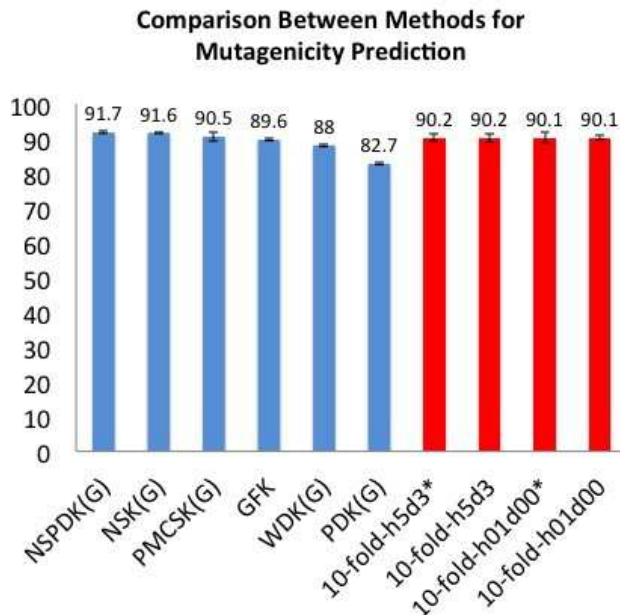


Figure 3. COMPARISON BETWEEN METHODS FOR MUTAGENICITY PREDICTION

random forest on un-kernelized height one and distance zero data did as well as or better than four previously published kernel-based methods (Figure 3). However random forest is not significantly better than any of the previously published methods (from two-sample unpooled t-tests). We also compared our random forest method against previously published methods for molecular interaction prediction from DrugBank data. The drug-target interactions with random forest had 94.5% AUROC and 86.1% validation accuracy (Figure 5) while a previously published kernel based method got 74% AUROC and 67.5% accuracy when they optimized for AUROC (Faulon et al., 2008). We cannot establish if our result is statistically significant because the other method did not report a standard deviation. However we believe 20% difference probably is statistically significant.

Discussion

Random forest on un-kernelized data does as well as or better than four previously published kernel-based methods in prediction mutagenicity. We hypothesize that this may be due to small local rules (such as “domain A with domain B is always mutagenic”) in chemistry that are not well approximated in the full dataset. Random forest uses subsets of training data and subsets of features to build the decision trees, which is more likely to capture these types of rules. The ran-

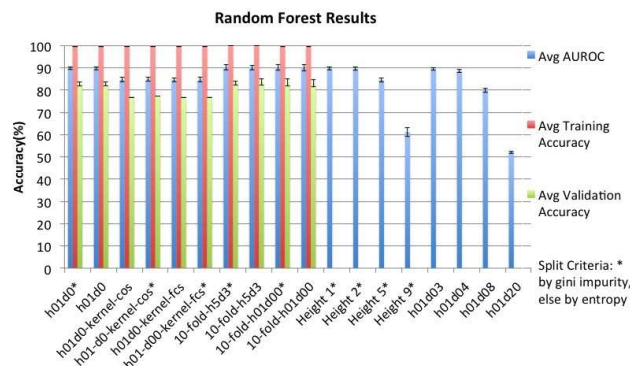


Figure 4. RANDOM FOREST RESULTS

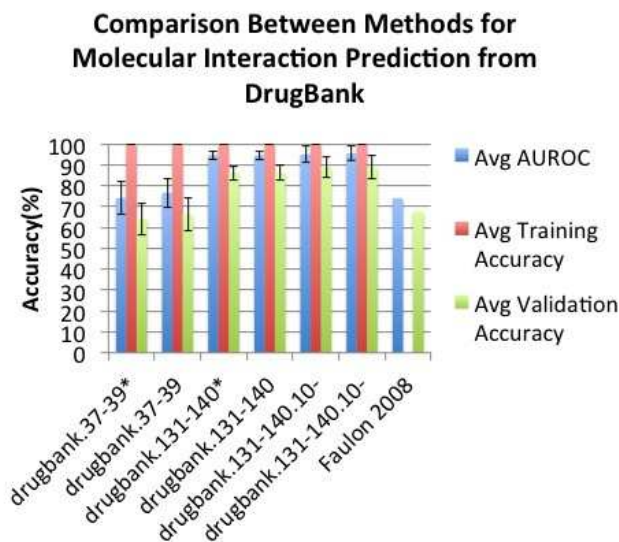


Figure 5. COMPARISON BETWEEN METHODS FOR MOLECULAR INTERACTION PREDICTION FROM DRUGBANK

dom forest implementation we used allows each decision tree to grow out fully, so possible features that can be used as splitting criteria depends on the subset of features used to generate each tree. This allows each tree to try to approximate potentially different small biological/chemistry rules. Height 1 and distance 0 produced the best results, which may seem counter-intuitive at first because the subgraphs themselves do not contain much information on the molecules topology. However, structure of individual trees may imply neighboring structures. Molecules are 3-dimensional structures, but we parsed and used their 2-dimensional structure so neighboring subgraphs in 3-dimensional space may not be captured with subgraph representation alone. Every splitting criteria in each tree of the

random forest is a feature representing the count of a small subgraph. Therefore, series of splits might represent a local structure found in the whole molecule. Another seemingly odd result we encountered is that using kernelized data did not help the AUROC or accuracies of our ensemble classifiers. The kernels we used were designed for use in SVMs and may not translate well into ensemble classifiers.

Drug-target interactions with random forest gets 94.5% AUROC and 86.1% validation accuracy while a previously published kernel based method got 74% AUROC and 67.5% accuracy. The other method used slightly different data that we tried to replicate as best as possible. We hypothesize that this may be due to small local rules (such as “domain A interacts with domain B) in chemistry that are not well approximated in the full dataset. Reasons outlined in the above as to why we believe random forest did better than previously published methods and why kernelizing the data did not help also applies here to molecular interaction prediction.

Related Work

Costa and Grave (2010) did not predict drug-target interactions, but used all feature heights up to a maximum height and all distances separating subgraph pairs up to a maximum distance (Costa & Grave, 2010). The authors achieved 91.7% AUROC for the graphs not augmented with functional groups from organic chemistry, improving to 92.0% AUROC with augmentation. Parsing all features for all small and large molecules in DrugBank was not computationally feasible in terms of time or memory.

Faulon et al. (2008) predicted drug-target interactions using support vector machines and a cosine similarity kernel, achieving 74% AUROC and 67.5% accuracy when optimizing for AUROC (Faulon et al., 2008). The authors used features at a single height for small molecules and a single height for proteins, rather than multiple heights.

Yu et al. (2012) also predicted drug-target interactions using random forest classifiers and achieved an average 89.65% AUROC (Yu et al., 2012). However, the feature set they used consisted of chemical and protein descriptors instead of subgraphs as in our project. Their dataset also is different than ours, but it is interesting to note that random forest has been applied to similar problem setups and has done well.

Future Work

1. We kernelized the data without considering interaction terms between features, so the next step is to kernelize the expanded feature set as well as trying to use the un-kernelized expanded feature set.
2. Try reducing the feature set based on results from a random forest before using other ensemble classifiers and kernels. Random forests can be used to assign variable importance to features. By using a reduced feature set, we might be able to more accurately capture small rules and potentially increase the accuracy of all the methods we tried in this project.
3. Use augmented graphs with annotated nodes and edges, using annotations such as domain family. Costa and Grave (2010) found a small improvement to their AUROC when using an augmented graph (Costa & Grave, 2010).
4. Take a larger sample from DrugBank. We are currently parsing 20,000+ positive training examples.
5. Continue developing ensemble classifier-based methods with kernels.

Conclusion

Random forest performed well in mutagenicity prediction and molecular interaction prediction compared to previously published methods. We show that random forest performs better than other ensemble classifiers in mutagenicity prediction and that parameter settings of height one, distance 0, and un-kernelized data produced the best results. Our random forest implementation did not do significantly better than previously published methods we found, but in terms of mean AUROC \pm one standard deviation, we did as well or better than 4 out of 6 other methods (Costa & Grave, 2010). Our random forest implementation for molecular interaction prediction did around 20% better in terms of AUROC and accuracy than a previously published method (Faulon et al., 2008).

Acknowledgements

Support was provided by the Tri-Institutional Training Program in Computational Biology and Medicine (via NIH training grant T32GM083937. The authors thank Professor Thorsten Joachims for teaching the class, and Adith Swaminathan for reviewing the project.

References

- Bolton, E., Wang, Y., Thiessen, P. A., and Bryant, S. H. PubChem: Integrated platform of small molecules and biological activities. *Annual Reports in Computational Chemistry*, 4(12):1–27, 2008. [ACS manuscript],[author manuscript].
- Costa, F. and Grave, K. D. Fast neighborhood subgraph pairwise distance kernel. In *Proceedings of the 27th International Conference on Machine Learning (ICML 2010)*, Haifa, Israel, 2010. [ICML manuscript].
- Faulon, J. L., Misra, M., Martin, S., Sale, K., and Sapra, R. Genome scale enzyme-metabolite and drug-target interaction predictions using the signature molecular descriptor. *Bioinformatics*, 2(50):225–233, 2008. PMID:18037612.
- Fukunishi, Y. Structure-based drug screening and ligand-based drug screening with machine learning. *Combinatorial chemistry & High Throughput Screening*, 12(4):397–408, 2009. PMID:19442067.
- Gaulton, A., Bellis, L. J., Bento, A. P., Chambers, J., Davies, M., Hersey, A., Light, Y., McGlinchey, S., Michalovich, D., Al-Lazikani, B., and Overington, J. P. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic Acids Research*, 40(Database issue):1100–1107, 2012. PMID:21948594.
- González-Porta, M., Frankish, A., Rung, J., Harrow, J., and Brazma, A. Transcriptome analysis of human tissues and cell lines reveals one dominant transcript per gene. *Genome Biology*, 14(R70):1–11, 2013. PMID:23815980.
- Kanehisa, M. and Goto, S. KEGG – kyoto encyclopedia of genes and genomes. *Nucleic Acids Research*, 28(1):27–30, 2000. PMID:10592173.
- Khandelwal, A., Krasowski, M.D., Reschly, E.J., Sinz, M.W., Swaan, P.W., and Ekins, S. Machine learning methods and docking for predicting human pregnane x receptor activation. *Chemical Research in Toxicology*, 21(7):1457–1467, 2008. PMID:18547065.
- Law, V., Knox, C., Djoumbou, Y., Jewison, T., Guo, A. C., Liu, Y., Maxiejewski, A., Arndt, D., Wilson, M., Neveu, V., Tang, A., Gabriel, G., Ly, C., Adamjee, S., Dame, Z. T., Han, B., Zhou, Y., and Wishart, D. S. Drugbank 4.0: shedding new light on drug metabolism. *Nucleic Acids Research*, 42(D):1091–1097, 2014. DOI:10.1093/nar/gkt1068, PMID:24203711.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. DOI:10.1109/5.726791,[author manuscript].
- Mitchell, Tom. M. *Machine Learning*. McGraw-Hill, 1997. ISBN 0070428077.
- Paul, S. M., Mytelka, D. S., Dunwiddie, C. T., Persinger, C. C., Munos, B. H., Lindborg, S. R., and Schacht, A. L. How to improve R&D productivity: the pharmaceutical industry’s grand challenge. *Nature Reviews Drug Discovery*, 9(9):203–214, 2010. PMID:20168317.

- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. [JMLR manuscript].
- Roy, A., McDonald, P. R., Sittampalam, S., and Chaguturu, R. Open access high throughput drug discovery in the public domain: a Mount Everest in the making. *Current Pharmaceutical Biotechnology*, 7(11): 764–778, 2010. PMID:20809896.
- Wishart, D. S., Jewison, T., Guo, A. C., Wilson, M., Knox, C., Liu, Y., Djoumbou, Y., Mandal, R., Aziat, F., Dong, E., Bouatra, S., Sinelnikov, I., Arndt, D., Xia, J., Liu, P., Yallou, F., Bjorn Dahl, T., Perez-Pineiro, R., Eisner, R., Allen, F., Neveu, V., Greiner, R., and Scalbert, A. HMDB 3.0 – The human metabolome database in 2013. *Nucleic Acids Research*, 41(D):801–807, 2013. PMID:23161693.
- Yu, H., Chen, J., Xu, X., Li, Y., Zhao, H., Fang, Y., Li, X., Zhou, W., Wang, W., and Wang, Y. A systematic prediction of multiple drug-target interactions from chemical, genomic, and pharmacological data. *PLoS ONE*, 7(5):e37608, 05 2012. doi: 10.1371/journal.pone.0037608. URL <http://dx.doi.org/10.1371/journal.pone.0037608>.